

Feb 2016

No.267

Religion and the Family: The Case of the Amish

James P. Choy

WORKING PAPER SERIES

Centre for Competitive Advantage in the Global Economy

Department of Economics

Religion and the Family: The Case of the Amish*

James P. Choy[†]

November 30, 2015

Abstract

I construct a model of religion as an institution that provides community enforcement of contracts within families. Family altruism implies that family members cannot commit to reporting broken contracts to the community, so the community must monitor contract performance as well as inflicting punishment. The community has less information than family members, and so community monitoring is inefficient. I provide evidence from a study of Amish institutions, including qualitative evidence from sociological accounts and quantitative evidence from a novel dataset covering nearly the entire Amish population of Holmes county, Ohio. I find that 1) Amish households are not unitary, 2) the Amish community helps to support families by inflicting punishments on wayward family members, 3) without the community Amish people have difficulty committing to punishing family members, and 4) Amish community membership strengthens family ties, while otherwise similar religious communities in which there is less need for exchange between family members have rules that weaken family ties. My model has implications for understanding selection into religious practice and the persistence of culture.

Keywords: Cultural Economics, Non-market Production, Public Goods, Religion

JEL Classification Numbers: D13, H4, Z10, Z12

*I am grateful to Robby Akerlof, Sascha Becker, Donald Kraybill, David McConnell, Steven Nolt, Debraj Ray, and various seminar participants for helpful comments. I thank David McConnell for data on Amish affiliations.

[†]Department of Economics, University of Warwick and CAGE. E-mail: j.choy@warwick.ac.uk

1 Introduction

One of the most important functions of religion is to regulate the family. From birth, to child raising, to marriage, to the care of elderly parents, religious teachings prescribe rules governing every aspect of family life. However, the literature on the economics of religion has so far ignored this family-regulating function. How do religious groups influence family relationships? Who chooses to join family-regulating religious groups? How can these groups persist over many generations? This paper tries to answer some of these questions.

I argue that religious institutions provide community enforcement of contracts within families. Contracts between family members differ from contracts between unrelated individuals because while family members may want to commit to contracts *ex ante*, family altruism implies that family members do not want to expose each other to punishment *ex post* by reporting broken contracts to the community. Thus the community must not only inflict punishment but must also monitor contract performance. However, the community has less information than the family about each family member's actions. As a result, community enforcement of family contracts is inefficient. I describe this inefficiency in the context of a society in which community enforcement of family contracts is particularly important, namely the Amish.

More specifically, I argue that the most important service provided by Amish communities is to help parents to control their children. I develop an argument based on Bernheim, Shleifer and Summers (1985). Parents would like their children to provide them with a non-monetary good, which I refer to as care. In return, they would like to bequeath money to their children. However, because parents are altruistic towards their children they cannot commit to withholding bequests, even from children who do not provide care. As a result, the level of care provided falls short of the parents' optimum.

The Amish community ameliorates this problem by imposing rules on all members of the community. Amish children are not allowed to attend school past the age of 14, and Amish adults are not allowed to use a wide variety of productivity enhancing technologies, including cars, electric tools and appliances, and many kinds of modern communication and information technology. Together, these rules reduce the wages of adult children. As a result, children who remain in the community spend less time working for a wage and more time caring for their parents. Parents incentivize children to remain in the community by committing to disinheriting children who leave. The reason parents can commit to disinheriting children who leave the community, but not to disinheriting children who merely fail to provide enough care, is that the community can observe whether children leave but cannot directly observe how much care each child provides. The community then exerts "peer pressure" on parents to disinherit children who leave, allowing parents to commit in a way that they would be unable to do on their own. Because the community and not the parents decide whether children should be punished, I say that Amish parents delegate the

monitoring of their children to the community.

A community that allows parents to commit to disinheriting children can help parents to control their children's actions. However, this mechanism is useful only for parents who are sufficiently rich. Poorer parents find that the threat of disinheritance, even when credible, is not sufficient to incentivize their children to remain in the community. To solve this problem, some Amish communities go beyond merely helping parents to commit individually. In some Amish communities, parents form a coalition where each parent agrees to inflict an additional punishment on any child who leaves by "shunning" such children. Since all the parents together can inflict a more severe punishment than any one parent individually, these communities can collectively incentivize children to remain even when each parent would be too poor to do so alone. In these communities, parents delegate not just the monitoring of their children but also part of the punishment of their children to the community.

In my model, Amish institutions are inefficient, in the sense that parents and children could achieve a Pareto superior outcome if parents could commit to a bequest schedule conditional their children's actions. The delegation of monitoring to the community creates the inefficiency in my model, because the community has less information about children's behavior than parents and so cannot impose optimal punishments for children who fail to provide enough care. The second-best monitoring scheme chosen by the community inflicts a deadweight loss by reducing children's wages, hurting both children and their altruistic parents.

I provide both qualitative and quantitative empirical evidence in favor of my model. For my quantitative evidence, I make use of the Ohio Amish Directory, a listing of nearly the entire Amish population of Holmes county, Ohio. I draw qualitative evidence from anthropological and sociological accounts of the Amish and a closely related group, the Hutterites. I have four findings. First, Amish households are not unitary, and there is conflict between parents and children. Some children want to leave the church, while their parents want them to stay. Second, Amish communities inflict punishments on children who leave the church, helping to support parents in this conflict. Third, without the support of the community Amish parents have difficulty in committing to punish their children. Fourth, Amish community membership strengthen family ties, while religious groups with similar theological beliefs and presumably similar preferences but in which parents have less need for care from their children have rules that weaken family ties. I argue that together, these findings provide strong support for my model.

My model is closely related to and inspired by the seminal club goods model of religion, developed by Iannacone (1992). However, there are some important differences. In the club goods model, religious groups impose rules to make people substitute away from private consumption and towards contributing to a religious club good that benefits the

entire community. In my model, the religious group also imposes rules that make children substitute away from private consumption, but the children substitute towards actions that benefit their parents rather than the whole community. This distinction has empirical consequences. Iannacone argues that religious groups described by the club goods model are most attractive to the poor. In contrast, in my model the purpose of the religious group is to allow parents to commit to using their resources in a wider variety of ways. This service is more useful to parents who have more resources, and so religious groups described by my model are most attractive to rich parents. Children of rich parents are also more likely to remain members of religious groups described by my model.

In reality many religious groups probably encourage group members both to take actions that benefit the whole community and to take actions that benefit their families. However, different religious groups differ in the extent to which they emphasize these different goals. Keeping this distinction in mind can help to reconcile seemingly contradictory empirical results. Iannacone finds that the poor are more likely to join strict Christian religious sects, while Botticini and Eckstein (2012) find that the rich were more likely to join ancient Jewish communities. Similarly, Chen (2010) finds that an exogenous decrease in income increases religious participation among Muslims in Indonesia, while Buser (2015) finds that an exogenous increase in income increases religious participation among Christians in Ecuador. Buser also finds that an exogenous increase in income makes people more likely to identify as evangelical Protestants rather than Catholics. I argue that these results indicate the existence of more than one kind of religious group. Right-wing religious groups emphasize “family values” and attract rich members, while left-wing religious groups emphasize community and charity and attract poor members.

The idea that some community members might form a coalition to prevent other community members from leaving also appears in other papers. Abramitzky (2008), Munshi and Rosenzweig (2006), and Munshi and Rosenzweig (2015) study communities where there is significant redistribution between member households.¹ In these communities richer households subsidize poorer households, and so rich households want to exit the community while poorer households want to force the richer households to stay. Poor households in these communities then form coalitions in various ways to increase the cost of leaving for rich households. In my model, there is also conflict in the community, but the conflict is not between rich and poor households but rather between parents and children within the same household. The coalition that forms is then not poor households trying to force rich households to stay, but rather parents from all households trying to force children from all households to stay.

My paper is also related to the community enforcement literature. A central concern in this paper is the information structure of the community, which determines how the commu-

¹Abramitzky studies the Israeli kibbutzim, while Munshi and Rosenzweig study the Indian caste system.

nity finds out when contract violations have occurred. Early contributions to this literature, such as Kandori (1992) and Greif (1993), assume that members of the community can costlessly observe when other members have broken contracts. More recent contributions, such as Bowen, Kreps and Skryzpacz (2013) and Ali and Miller (2015) assume instead that the parties to a contract must report to the community whether the contract has been fulfilled. These papers construct equilibria in which each party to a contract is just indifferent between reporting that the contract has been fulfilled and that the contract has been broken, so that it is incentive compatible for each party to report the outcome of the contract truthfully. I assume that the community cannot observe the outcomes of the contracts between families, and also that because of family altruism family members always strictly prefer not to report broken contracts to the community. This problem necessitates the inefficient community monitoring scheme in my model.

Finally, my paper is related to theories of cultural evolution developed by Boyd and Richerson (1985) and Bisin and Verdier (2000, 2001). Both of these theories attempt to explain the existence of cultural behaviors that are seemingly welfare reducing relative to some apparently feasible alternative. Both theories deny that these seemingly welfare reducing cultural behaviors are really welfare reducing. Boyd and Richerson argue that cultural behaviors are “rules of thumb” for boundedly rational agents. Conditional on bounded rationality, agents maximize their welfare by applying these rules when they cannot discover the unconditionally optimal action. Bisin and Verdier argue that seemingly welfare reducing cultural behaviors are in fact optimal responses by agents who have unusual preferences. A similar argument appears in Akerlof and Kranton (2000). In contrast to these theories, I argue that Amish culture really is welfare reducing relative to a world in which Amish parents and children could enter into complete contracts.

In addition to explaining why seemingly welfare reducing cultural behaviors exist, my theory has implications for understanding when and how minority cultural traits can persist across generations. There are two main theories of cultural persistence in the economics literature. Lazear (1999) and Michalopoulos (2012) argue that it is difficult for members of minority cultural groups to learn the majority language, and that this technological barrier prevents minority groups from assimilating into the majority. Bisin and Verdier (2000, 2001) argue that parents in minority groups prefer that their children share the parents’ cultural traits, and that the parents also pass this preference on to their children. These preferences then sustain the distinctiveness of the minority group. In contrast to these theories, I argue that neither language barriers nor parental preferences on their own are sufficient to sustain Amish cultural distinctiveness in the face of material incentives favoring assimilation. However, a parental coalition dedicated to collective punishment of children who assimilate into the majority culture can be very effective in preventing assimilation over multiple generations.

The remainder of this paper is as follows. Section II provides background on the Amish. Section III presents the model. Section IV presents both qualitative and quantitative empirical evidence regarding the Amish. Section V provides further evidence from other religious groups. Section VI concludes.

2 Background on the Amish

The Amish are one of several Protestant sects, collectively known as Anabaptists, which emerged in Switzerland and Germany in the aftermath of the Protestant Reformation.² The Anabaptists were among the most radical Protestant groups, in the sense that they went the furthest in rejecting the traditional authority structures of medieval life. In particular the Anabaptists rejected the traditional church hierarchy, instead holding meetings in private homes with lay preachers, and they also denied the right of the ruler of a territory to determine the religion of the territory's inhabitants. The Anabaptists believed instead that each person should make an informed choice of what religion to belong to on reaching adulthood. One consequence of this belief was that the Anabaptists did not practice infant baptism, hence their name. The Amish denomination was founded in 1693 in Switzerland by Jakob Amman, who split from the existing Anabaptist congregations over the issue of shunning. I discuss shunning in more detail below.

The Amish faced persecution in Europe due to their religious beliefs, and as a result many Amish people emigrated to the United States and Canada in the 18th and 19th centuries. The remaining Amish church in Europe did not prosper, and the last European Amish congregation disbanded in 1937. In the Americas, however, the Amish church grew quickly. Table 1 shows Amish population estimates from 1900 to the present. This growth is due almost entirely to the high Amish fertility rate and not to conversion into the Amish church from outside, which is nearly impossible. In fact the Amish population has long been a subject of genetic research due to its extreme genetic isolation.

Amish theology emphasizes humility, simplicity, sublimation of personal desire, and surrender to God's will. One aspect of this theology is separation from the outside world, which the Amish regard as materialistic and corrupt. As part of their separation from the outside world, the Amish speak a dialect of German called Pennsylvania German, which is related to the dialect of southern Germany where the Amish faith originated. Amish children are also taught to read standard German, and Amish church services are held in standard German using Martin Luther's 1534 German translation of the Bible, although fluency in standard German varies widely. Most Amish people are also fluent in English and speak English with outsiders.

²This overview of Amish history, culture, and institutions is drawn from Beachy (2011), Hostetler (1993), Hurst and McConnell (2010), Kraybill (2001), and Kraybill, Johnson-Weiner, and Nolt (2013).

2.1 Amish community organization

Contemporary Amish society is organized into settlements, affiliations, and districts. A settlement is a geographical area with many Amish people. The largest settlement, centered in Holmes county, Ohio, and including parts of six surrounding counties, has more than 30,000 members. There are two other settlements of similar size, in Lancaster county, Pennsylvania, and Elkhart and Lagrange counties, Indiana, and many smaller settlements. When these settlements were originally formed they were located in rural areas, and traditionally all Amish men were farmers. However, in recent years urban development has encroached on Amish settlement areas and now many parts of the larger Amish settlements are more suburban or exurban than rural. Urban expansion and rising land prices have led many Amish people to abandon farming for other trades. In the Holmes county settlement, which is the main object of study in this paper, many Amish men work as artisanal craftsmen in areas such as woodworking and masonry, or as small business owners in businesses such as retail or construction. Other Amish men work as unskilled or semi-skilled laborers, including for non-Amish employers. Finally, some Amish men remain in the traditional occupation of farming. Married women rarely work outside the home. Settlements do not have well-defined borders, and within each settlement Amish people and non-Amish people live side by side.

Except in the Lancaster county, Pennsylvania settlement, there are no settlement-wide governance institutions. Governance instead happens at the level of the district. A district is a collection of 20-40 households that meets once every two weeks in a member's house or barn to hold religious services. The size of a district is limited by the number of people who can fit in a house or barn for the service, and when a district becomes too large, it splits. Each district is responsible for making and enforcing its own rules, which I describe in more detail below. Rules must be agreed upon unanimously by all adult members of the church at special meetings held twice each year. While in principle all members have an equal say in these meetings, in practice the district leader, known as the bishop, has significantly more power than other district members through his ability to set the agenda. Bishops are chosen for life from among the male elders of each district through a complex procedure that is partly an election but that also involves elements of chance. All adult district members, both male and female, have a vote in these elections.

Amish society features a substantial amount of mutual insurance and redistribution, and much of this insurance is organized at the level of the district. For example, each district has an officer called the deacon, whose responsibilities include managing an alms fund for less fortunate district members. Wealthier district members may also be required to provide no-interest loans for poorer members to buy houses or start businesses. At the same time, however, Amish church members are typically required to refrain from taking advantage of insurance and redistributive institutions provided by the larger society. The Amish

are specifically exempted from mandatory participation in Social Security and government-provided health insurance, and usually do not accept workplace-provided health insurance or government-provided aid to the poor such as food stamps. Overall it is unclear whether Amish insurance and redistributive institutions provide more or less insurance and redistribution than the institutions of the larger society that they replace.

While each district is responsible for deciding its own rules, in most cases groups of districts choose similar rules. A group of districts with similar rules is called an affiliation, and two districts in the same affiliation are said to be “in fellowship”. Districts in fellowship may exchange guest preachers for Sunday services. In the Holmes county, Ohio settlement, there are three major affiliations, the Old Order, the New Order, and the Andy Weaver affiliation. (A fourth major affiliation, the Swartzentrubers, are not included in my data set and so are not considered in this study.) The Old Order are the largest affiliation and represent the main line of the Amish tradition, going back to the Amish settlement of Ohio in the 19th century. The Andy Weaver affiliation split from the Old Order in 1954 and the New Order split from the Old Order in 1958.

2.2 Amish rules

Amish districts impose many rules on their members. Roughly speaking, these rules fall into four categories. First, there are rules about the use of technology. Second, there are rules about family life and child raising. Third, there are rules about dress and appearance. Fourth, there are rules about the punishment of district members who break the other rules of the district.

Amish rules about the use of technology are the aspect of Amish culture that is perhaps best known to outsiders. The rules forbid the use of a wide variety of modern technologies. Most notably, all districts prohibit driving cars. Instead Amish people travel in horse-drawn buggies, and these buggies are a common sight on roads in and near Amish settlements. Many districts also have rules about the use of other kinds of transportation and motorized equipment. For example, many districts have rules about the use of tractors. Some districts permit all tractors, some districts permit only tractors with steel wheels, which cannot be driven on the road, and some districts forbid tractors entirely. Some districts also prohibit bicycles. Most Amish districts prohibit the use of electricity from the grid, although some districts permit electric appliances powered by gas generators. Finally, most districts have rules against modern communication and information technology such as telephones and computers, although the extent to which these technologies are prohibited varies from district to district.

Amish districts have many rules and customs regarding child raising and family life, including rules about proper gender roles in marriage, courtship customs, and child discipline. From the perspective of an economist the most important of these rules is the rule

prohibiting children from attending school past the age of 14. This rule is followed in all districts, and as a result all Amish people have the same amount of education.³ The Amish were specifically exempted from mandatory state education laws in the 1972 Supreme Court case *Wisconsin v. Yoder*.

Rules about dress and appearance serve to distinguish the Amish visually from non-Amish people. Men and women both dress modestly, and men have beards but not moustaches and distinctive haircuts. There are subtle differences between the clothing required by different affiliations, such as different rules across affiliations concerning the width of brims for men’s hats. For someone familiar with the culture it is possible to distinguish members of different affiliations in this way.

Districts enforce the rules using a penalty called shunning. A shunned member is not allowed to interact in certain ways with church members in good standing. For example, it is forbidden to share a meal with or to accept gifts from someone who has been shunned, although interestingly it is still permitted to give gifts to a shunned member. Shunning may be temporary or permanent. Temporary shunning lasts for two to six weeks and is a response to less serious violations of the rules. A church member who repeatedly breaks the rules, or who refuses to show remorse for past rule-breaking, may be permanently shunned. Permanent shunning lasts indefinitely, but even a member who has been permanently shunned may be accepted back into the community after making a confession and repenting in front of the congregation. Information about shunning is shared across districts in the same affiliation, and if it is discovered that a person has been shunned in his or her home district, members of other districts in the same affiliation will respect the prohibition against interacting with that person.

As I discuss below, in some cases church members may be shunned for voluntarily leaving the church, even if they have not otherwise done anything wrong. Church members who leave but who are not shunned may continue to interact with their friends and family who remain Amish, and in fact such continued interaction is common. Continued interaction is possible because leaving the church does not necessarily imply moving geographically. For example, Amish parents may rely on their children who have left the church to run errands that require access to a car.

2.3 Variation in rules across districts

The Amish describe the variation in rules and ideology across affiliations as being on a scale from “low” to “high”. Lower affiliations are more removed from the outside world, while higher affiliations interact more with non-Amish society. Of the three affiliations that I study in this paper, the Andy Weaver affiliation is the lowest, the Old Order is in the

³Kraybill (2001, p. 81) reports the results of a survey of educational attainment among 812 Amish adults in Lancaster county, Pennsylvania. He finds zero who have more than an eighth grade education.

middle, and the New Order is the highest.

Districts in low affiliations typically are more restrictive about technology use than high affiliations. For example, districts in lower affiliations may forbid bicycles, tractors, and gas-powered appliances, while districts in higher affiliations may permit the use of these technologies. Interestingly, the scale is reversed when it comes to rules about family life and child-rearing practices. Parents in lower affiliations are somewhat more lenient than parents in higher affiliations in allowing their children to hold unsupervised social events, and lower affiliations are more accepting of alcohol and tobacco use than higher affiliations.

Although there are some differences between affiliations in the rules governing technology use and family life, these differences are not the most important differences between affiliations. Instead, the central point of disagreement between affiliations concerns the proper use of shunning. Districts that have similar attitudes towards shunning are likely to be in fellowship even if they have divergent attitudes towards technology or family life. For example, the Lancaster county, Pennsylvania Amish districts are in fellowship with the Andy Weaver affiliation in Ohio, because they share the same policies regarding shunning, even though the Lancaster county Amish permit the use of more modern technologies than the Andy Weaver affiliation.

Affiliations disagree in particular about the extent to which shunning should be applied to members who have voluntarily left the church but who have not otherwise done anything wrong. The Andy Weaver affiliation practices strict shunning or *streng meidung*. This means that any member who leaves the church is permanently shunned. In principle this penalty applies even to members who leave to join a higher Amish affiliation, although in practice there is some leeway for members who leave the Andy Weaver affiliation but remain Amish. Shunning can be ended only if the wayward member returns to an Andy Weaver district. In the Old Order, shunning practices vary from district to district and from case to case. Members who leave to join other Amish affiliations are rarely shunned. Members can also usually leave to join other Anabaptist denominations without incurring a penalty. This typically means joining the Mennonites, an Anabaptist group that is theologically very similar to the Amish but which does not impose nearly as many rules on its members. The area of the Holmes county settlement also includes many Mennonites, and so there is an alternative community available for people who leave the Amish church. Members who leave to join other Christian denominations, or who become atheists, are more likely to be shunned by their former Old Order districts. New Order districts do not shun former members as long as those members join a “Bible-believing church”, and many New Order Amish believe that the practice of shunning should be abandoned completely. In principle, it is possible for children to leave any affiliation without incurring a penalty if they leave before being baptised, where baptism typically happens in the late teens or early twenties. In practice, however, there is considerable social pressure on children to be baptised. Community members may

also “draw back” from interaction with children who leave the church before being baptised, even if the children are not officially shunned. Like the official punishment of shunning, this unofficial punishment is more likely to be applied by lower affiliations.

3 Theory

In this section I develop a model to explain the unique features of Amish society described above. In particular I want to explain why the Amish church imposes productivity reducing rules, such as rules limiting children’s education and rules against using modern technologies, and why some Amish districts shun former members who have left the church. I argue that these features of Amish society allow parents to control their children in ways that would be impossible in the absence of the rules.

The model is as follows. Society consists of parents and children. Each parent (she) has one child (he). Each child has one unit of time, which can be allocated either to working for a wage w or to caring for his parent. Each child’s utility is

$$u_c(x_c, t)$$

where x_c is the child’s consumption and t is the time that the child spends caring for the parent. Parents are altruistic towards their children. Each parent gets utility from her own consumption, from time that the child spends caring for her, and also from her child’s utility. Each parent’s utility function is

$$u_p(x_p, t, u_c(x_c, t))$$

where x_p is the parent’s consumption. I assume that both the parent’s and the child’s utility functions are increasing in all of their arguments. Each parent is endowed with wealth y which she can either consume or bequeath to the child. Let b be the amount of the bequest. The parent’s budget constraints are

$$\begin{aligned} x_p + b &\leq y \\ b &\geq 0 \end{aligned}$$

The child’s budget constraints are

$$\begin{aligned} x_c &\leq w(1 - t) + b \\ 0 &\leq t \leq 1 \end{aligned}$$

The child first chooses his allocation of time between working and caring for the parent, and the parent then chooses her bequest.

After observing the child's choice of t and the child's wage w , the parent effectively chooses her own and her child's consumption subject to the constraint that bequests must be non-negative and that total consumption must be less than or equal to the sum of the resources controlled by the parent and the child. The parent's optimization problem can be written as

$$\max_{x_p, x_c} u_p(x_p, t, u_c(x_c, t))$$

subject to the constraints

$$\begin{aligned} x_p + x_c &\leq y + (1 - t)w \\ x_c &\geq (1 - t)w. \end{aligned}$$

Let $x_c(t, w)$ be the solution to this problem for each choice of t by the child. I assume that x_c is a normal good for the parent, that is, that $x_c(t, w)$ is weakly increasing in w and y for all t . The child's problem is then to choose a point on the $x_c(t, w)$ curve to maximize his utility. Figure 1 displays this problem graphically. The kink in the $x_c(t, w)$ curve represents the level of t below which the parent does not give the child a bequest. As can be seen in the figure, as w decreases the child's equilibrium utility falls. Since $x_c(t, w)$ is also increasing in y , the same argument shows that as y decreases the child's equilibrium utility falls.

In contrast, an decrease in w may increase the parent's equilibrium utility. To show this, I take the derivative of the parent's utility function with respect to w to get:

$$\frac{du_p}{dw} = \frac{\partial u_p}{\partial x_p} \frac{\partial x_p}{\partial w} + \frac{\partial u_p}{\partial t} \frac{\partial t}{\partial w} + \frac{\partial u_p}{\partial u_c} \frac{\partial u_c}{\partial x_c} \frac{\partial x_c}{\partial w} + \frac{\partial u_p}{\partial u_c} \frac{\partial u_c}{\partial t} \frac{\partial t}{\partial w} \quad (1)$$

In analyzing this expression, it is helpful to consider separately the case in which the parent gives a positive bequest and the case where the parent gives no bequest. Suppose first that the parent gives no bequest. In this case $\partial x_p / \partial w = 0$, so equation (1) simplifies to

$$\frac{du_p}{dw} = \frac{\partial u_p}{\partial t} \frac{\partial t}{\partial w} + \frac{\partial u_p}{\partial u_c} \frac{du_c}{dw} \quad (2)$$

The second term on the right hand side of (2) is positive but the first term may be negative, if $\partial t / \partial w$ is negative. If the first term is sufficiently negative then du_p / dw is negative. Intuitively, increasing w causes the child to substitute away from caring for the parent and towards working for a wage. This increases the child's utility, which increases the parent's utility, but the decrease in t lowers the parent's utility. If the decrease in the

parent's utility from the decrease in t is sufficiently large, then the overall effect on the parent's utility may be negative.

Now consider the case in which the parent gives the child a positive bequest. In this case, noting that $x_c + x_p = y + (1 - t)w$ and therefore that $\partial x_c / \partial w + \partial x_p / \partial w = (1 - t)$, and also that $\partial u_p / \partial x_p = (\partial u_p / \partial u_c)(\partial u_c / \partial x_c)$ since the parent chooses her bequest optimally, equation (1) simplifies to

$$\frac{du_p}{dw} = (1 - t) \frac{\partial u_p}{\partial x_p} + \frac{\partial t}{\partial w} \left(\frac{\partial u_p}{\partial t} + \frac{\partial u_p}{\partial u_c} \frac{\partial u_c}{\partial t} \right) \quad (3)$$

When the parent makes a positive bequest to the child, while holding the child's time allocation constant an increase in the child's wage has the same effect as a direct increase in the parent's consumption, since the parent can trade off her own consumption one-for-one with the child's consumption. The first term of (3), which is positive, represents this effect of an increase in the child's wage on the parent's utility. The second term represents the effect of a change in t due to a change in the child's wage on the parent's utility. As before, if $\frac{\partial t}{\partial w}$ is negative, then this second term is negative, and if $\frac{\partial t}{\partial w}$ is sufficiently negative, then du_p/dw may be negative.

Whether or not the parent leaves the child a bequest, then, it is possible that a decrease in w increases the parent's utility. Proposition 1 summarizes the previous discussion:

Proposition 1 *If the child's consumption x_c is a normal good for the parent, then*

1. *A decrease in w always weakly reduces the child's equilibrium utility.*
2. *A decrease in y always weakly reduces the child's equilibrium utility.*
3. *A decrease in w may increase the parent's equilibrium utility.*

So far my analysis is very similar to the analysis in Bernheim, Shleifer and Summers (1985).⁴ I now introduce a new element by modifying the game to consider the role of the Amish community in helping the parent to impose her wishes on the child. Suppose that the child can either become a member of the Amish church or not. If the child is not a member of the church, then the child receives wage \bar{w} . The church imposes restrictions on the child's education and on the productive technologies that the child can use, so that if the child becomes a member of the church then the child receives wage \underline{w} , where $\underline{w} < \bar{w}$. Let m denote the child's membership in the church, where $m = 1$ if the child is a member of the church and $m = 0$ otherwise. In addition, the parent can partially commit to a bequest schedule that depends on whether the child is a member of the church or not. In particular, the parent can commit to leaving the child a bequest $b(m) = 0$ if $m = 0$. However, if the child does join the church the parent cannot commit to any bequest.

⁴See also Ghosh and Karaivanov (2007).

The idea here is that the Amish community can observe whether the child is a member of the church, and whether the child follows the rules of the church if he is a member, but cannot directly observe how much care the child provides to the parent. If the child leaves the church, the Amish community exerts “peer pressure” on the parent to withhold a bequest from the child, as in Kandel and Lazear (1992). Thus the parent can commit to disinheriting a child who leaves the church. However, the community cannot bring its influence to bear on the parent if the child merely fails to provide enough care. As a result the parent cannot commit to a bequest schedule conditional the amount of care that the child provides.

The timing of the new game is as follows:

1. Parent chooses whether to commit to $b(0) = 0$.
2. Child chooses whether to join the church.
3. Child chooses t .
4. If the parent did not commit to $b(0) = 0$ in step 1, or if the child chose to join the church in step 2, the parent chooses a bequest $b \geq 0$. Otherwise the parent gives the child no bequest.

Now consider how the parent’s wealth affects the child’s decision to join the church. Let $v_p(w, y)$ be the equilibrium utility of a parent who has not committed to $b(0) = 0$, who has wealth y , and whose child receives wage w . I make the following assumption:

Assumption 1

For all y ,

$$v_p(\underline{w}, y) > v(\bar{w}, y)$$

Assumption 1 states that each parent prefers that her child join the church, even though doing so lowers the child’s wage. Assumption 1 implies that children who join the church provide more care to their parents than children who do not, since the only reason that the parent might want the child to join the church is to induce the child to provide more care. Let $v_c(w, y)$ be the equilibrium utility of a child whose parent has not committed to $b(0) = 0$, who receives wage w , and whose parent has wealth y . If the child’s consumption is a normal good for the parent, then from proposition 1, $v_c(w, y)$ is increasing in both y and w . If the child joins the church, then he gets utility $v_c(\underline{w}, y)$, while if he does not join the group then he gets utility $v_c(\bar{w}, 0)$. For sufficiently large y , $v_c(\underline{w}, y) > v_c(\bar{w}, 0)$, and so if the parent commits to disinheriting the child if the child does not join the church, then the child prefers to join the church.

Proposition 2 summarizes the previous discussion:

Proposition 2 *Suppose that the child's consumption is a normal good for the parent and that assumption 1 holds. Then there exists \bar{y} such that the following strategies are part of a subgame perfect Nash equilibrium:*

1. *If and only if $y < \bar{y}$, the parent does not commit to disinheriting the child if the child does not join the church, and the child does not join the church.*
2. *If and only if $y \geq \bar{y}$, the parent commits to disinheriting the child if the child does not join the church, and the child joins the church.*

Proposition 2 states that children of rich parents join the church while children of poor parents do not. The reason is that the parent's threat to disinherit the child if he does not join the church is effective only if the parent is sufficiently rich.

Now, suppose that the community can inflict an additional punishment p on a child who does not join the group. Then the utility of a child who does not join the group is $v_c(\bar{w}, 0) - p$. We have the following:

Proposition 3 *The cutoff parental wealth \bar{y} below which a child does not join the group is decreasing in p .*

The punishment p in the model corresponds to the real-world probability that the child will be shunned if he leaves the Amish church. Thus p is larger in lower Amish affiliations, and so my model predicts that children are less likely to exit from lower affiliations.

I make two remarks on the properties of the equilibria described above:

Remark 1 *The equilibria I describe are Pareto inefficient.*

The wage penalty for being a member of the Amish church is a deadweight loss. Both the parent and the child would be better off if the child could leave the church and get a higher wage while maintaining the same bequest and care levels. However, because the parent cannot commit either to disinheriting the child for failing to provide enough care, or to reporting the child's misbehavior to the community for punishment by the community, this Pareto improving allocation cannot be achieved.

Remark 2 *Religious group membership is most attractive to rich parents.*

The religious group has no effect on the child's behavior if the parent has wealth less than the cutoff \bar{y} . Thus for poor parents, religious group membership has no value. In contrast, the religious group causes the child to provide more care to the parent if the parent has wealth greater than the cutoff, and so religious group membership may increase the utility of rich parents. Thus, not only are children of rich parents more likely to remain members of the religious group, but rich parents are also more likely to join the group if they have the opportunity to do so.

3.1 Comparison to the club goods model

I briefly compare the predictions of my model to the club goods model. For a detailed presentation of the club goods model, see Iannacone (1992).

In the club goods model, religious group members can allocate resources either to private consumption or to contributions to a religious club good that benefits all members of the community. Because contributions to the club good yield positive externalities for other members of the community, contributions to the club good are undersupplied in Nash equilibrium relative to the social optimum. The religious group imposes rules prohibiting certain kinds of private consumption in order to make people substitute away from private consumption and towards contributions to the club good. Iannacone shows that under some circumstances, these rules may be welfare improving for all members of the community.

In the context of the club goods model, altruistic expenditures of time or money to benefit family members are just another kind of private consumption. Thus, religious groups described by the club goods model may be neutral towards the family, or they may even make rules that decrease the quantity of resources that group members devote to their families in order to encourage group members to contribute those resources to the community instead. This implication of the club goods model sharply contrasts with my model, in which the religious group makes rules designed to increase the quantity of resources that group members devote to their families.

Iannacone is agnostic about exactly what religious group members contribute to the community. As an example, he suggests enthusiastic participation in religious services. However, the following literature has argued that the most important religious club good is charitable contributions of time or money to aid poorer or less fortunate members of the community. This interpretation of the club good as mutual insurance or redistribution appears in Berman (2000), Dehejia, DeLeire, and Luttmer (2007), and Chen (2010).

Religious groups that provide significant amounts of aid or redistribution are likely to be most attractive to the poor. In contrast, as discussed above religious groups described by my model are likely to be most attractive to rich parents, and children of rich parents are also more likely to remain members of these groups. This is a second important distinction between my model and the club goods model.

The previous discussion suggests an empirical correlation. Religious groups described by my model feature relatively little redistribution, attract rich members, and have rules that increase the quantity of resources that group members devote to their families. Religious groups described by the club goods model feature significant redistribution, attract poor members, and have rules that are neutral towards the family or that reduce the quantity of resources that group members devote to their families. Thus I make the following prediction, which for consistency with the previous results I state as a “proposition”:

Proposition 4 *There is a correlation between the amount of redistribution in a religious*

group, the wealth of the religious group's members, and the religious group's attitude towards the family. Religious groups with relatively less redistribution have rules that increase the quantity of resources that family members devote to each other and that strengthen family ties, and these groups attract wealthier members. Religious groups with relatively more redistribution may have rules that decrease the quantity of resources that family members devote to each other and that weaken family ties, and these groups attract poorer members.

In the next section I test this hypothesis together with the implications of propositions 1 through 3.

4 Empirics

In this section I provide empirical evidence that the mechanism described above explains Amish institutions. The argument proceeds in four steps. First, I argue that Amish households are not unitary. In particular, at least some Amish children prefer to leave the church, while their parents want them to stay. Second, I argue that Amish communities support parents in this conflict by helping to punish children who leave the church. Third, I argue that without support from the community, parents have difficulty committing to punishing their children. Finally, I argue that Amish parents want their children to join the church because church membership strengthens family ties. I argue that together, these four claims establish that the mechanism described above is at least part of the explanation for Amish rules. Before proceeding with the argument, however, I describe my data sources.

4.1 Data

I combine data from two sources. The first source is the Amish directory for the Holmes county settlement. The Amish directory is a list of all the households in the settlement, with the exception of households in the Swartzentruber affiliation who do not participate. It is published irregularly every 5-10 years. For each household, it contains the names and birth dates of all members of the household, the profession of the household head, and the names of all of the children of the household head and the spouse of the household head. For each child the directory states whether the child is a member of the Amish church and whether the child lives in the settlement. The directory also gives addresses for each child. The main outcome variables that I am interested in are whether a child leaves the church and whether a child migrates geographically. I consider a child to have migrated if the child is a member of the Amish church living in a different settlement, or if the child is not a member of the Amish church and has an address in a state other than Ohio. It is important to emphasize that leaving the church and migrating are independent decisions, and I observe children who have left the church but not migrated who have migrated but not left the church, both, and

neither.

I make use of two editions of the Amish directory, the 1988 edition and the most recent 2010 edition. My sample is all children between the ages of 8 and 16 who are living in the settlement in 1988. I match this sample with the 2010 directory in order to find outcomes for children in 2010, when the children are between the ages of 30 and 38. Since decisions to leave the church are mostly made early in life, it is likely that most of these children have permanently chosen their religious status by age 30. I observe children's outcomes through entries for their parents in the 2010 Amish directory, so I am able to observe outcomes even for children who have left the church or who have migrated, as long as their parents remain members of the Holmes county Amish church. I do not observe outcomes for children of parents who die between 1988 and 2010, or for children of parents who leave the church or the settlement between 1988 and 2010. I am able to observe outcomes for 93% of the children in my sample.

I match this data with 1988 property tax data collected from the county treasurer's offices of the counties that contain the Holmes county settlement. Not surprisingly the majority of the children in my sample live in Holmes county, but the settlement also includes parts of neighboring Tuscarawas, Coshocton, Wayne, and Stark counties, Ohio. I observe the assessed value of buildings and land owned by each parent in 1988. Table 2 shows summary statistics.

The theoretical discussion above suggests that the wealth distribution of a religious group is important for understanding the group's economic function. Figure 2 shows a histogram of the value of buildings owned by each household in my sample. There are two important takeaways from figure 2. First, Amish households are relatively wealthy. The median value of buildings owned by all households is \$60,729, and the median value of buildings owned by households who own any buildings is \$66,200. The median value of all homes in Ohio in 1988 was \$60,957, so the Amish are as wealthy or perhaps slightly wealthier than other Ohio households. This is remarkable given Amish restrictions on schooling and technology. I hypothesize that Amish culture has non-cognitive benefits for children that compensate for Amish educational deficits. Second, there is significant wealth inequality across Amish households. There does not appear to be dramatically more redistribution among the Amish than in the larger society.

I now proceed to argue that the data support my theory of Amish rules.

4.2 Step 1: Amish households are not unitary

The first step of my argument is to show that children of richer Amish parents are more likely to remain in the church, consistent with proposition 2 above. I argue that this result implies that Amish households are not unitary, and that some Amish parents want their children to remain in the church while their children prefer to leave.

To show that children of richer parents are more likely to remain in the church, I run regressions of the form

$$Y_{ihd} = \alpha + \beta_1 P_{hd} + \beta_2 \mathbf{X}_{ihd} + \epsilon_{ihd}$$

Here $Y_{ihd} = 1$ if child i in household h in district d has left the church by 2010, and $Y_{ihd} = 0$ otherwise. P_{hd} is a measure of the permanent income of the parents in household h in district d in 1988, \mathbf{X}_{ihd} is a vector of control variables, and ϵ_{ihd} is an error term. I cluster error terms by district.

My preferred measure of the permanent income of household h is the logarithm of the total value of buildings owned by the parents in household h in 1988. The value of buildings owned by the household is a measure of the flow value of housing services consumed by the household, which in turn is a (rough) proxy for the household's permanent income. It is important to note that the value of buildings owned by a household is not a proxy for the value of that household's asset holdings, since I do not observe the value of any mortgages owed.

I also observe the value of land owned by each household. However using the value of land owned by each household as a proxy for permanent income is problematic, for two reasons. First, many households own significant amounts of agricultural land, which is a productive asset rather than a durable consumption good, and so including land will create an upward bias in the estimate of each household's permanent consumption. Second, there are two regulatory regimes for assessing land values for tax purposes in Ohio. Some parcels of land are assessed at market value, while other parcels are assessed according to their Current Agricultural Use Value (CAUV). CAUV values are lower than market values by up to a factor of ten, and there is no easy way to convert CAUV values to market values. For most households, I observe both the market value and the CAUV value of land. However, for households in Wayne and Tuscarawas counties, I observe only the CAUV value for parcels that fall under the CAUV regime. Thus, values of CAUV parcels in Wayne and Tuscarawas counties are not comparable to parcels in other counties or to parcels in the same county that do not fall under the CAUV regime. Despite these problems, as a robustness check I also run regressions using the logarithm of the combined value of land and buildings owned by a household as a measure of that household's permanent income. In these regressions, I drop observations from Wayne and Tuscarawas counties.

A significant number of households do not own any property. These households are presumably not homeless, but rather are either renting housing, or (more likely in this context) living with other family members. Thus households who own no property are not properly interpreted as households whose flow value of housing consumption is zero, but rather as households whose flow value of housing is unobserved. Following this interpretation, in most regressions I drop households who own no property.

Table 3 reports results. Column 1 shows a regression with no controls other than affiliation fixed effects. Children of richer parents are significantly less likely to leave the church, consistent with proposition 2 above. Doubling the value of buildings owned by a household reduces the probability that a child from that household leaves the church by 5 percentage points, which is a large effect given the mean probability of leaving the church of 14.7%.

Column 2 adds control variables. The coefficient on parents' wealth is similar to the coefficient in the regression without controls. Children of farmers are significantly less likely to leave the church than children of non-farmers. This is consistent with the fact that farmers bequeath a larger fraction of their wealth to their children, as shown, for example, by Behrman and Rosenzweig (2006). Thus for a given level of parental wealth the threat of disinheritance is more severe for children of farmers than for children of non-farmers. Boys are significantly more likely to leave the church than girls. This is likely due to the fact that boys learn skills that make them employable in the market economy, while girls mainly learn homemaking skills that are less useful outside of the Amish community. I also control for each child's age and the ages of the child's mother and father.

One possible problem with the regression in column 2 is that different households may spend different proportions of their total permanent income on housing, and these differences may be correlated with other household characteristics that determine children's propensity to exit the church. In this case using housing consumption as a measure of permanent income generates biased estimates of the effect of parents' permanent income on children's propensity to exit the church. In order to partially control for this possibility, I include each household's completed fertility in 2010 as a control variable in column 3. Households with higher fertility are likely to spend a larger fraction of their permanent income on housing. Reassuringly, including this control does not meaningfully change the estimate of the effect of parents' wealth on the probability of exit.

Column 4 repeats the regression in column 2 using the logarithm of the total value of buildings and land owned by each household as a measure of that household's wealth. As discussed above, in this regression I drop observations from Wayne and Tuscarawas counties. The results are qualitatively similar to the results in column 2.

Column 5 regresses the probability of exit on an indicator for whether the child's parents own any property. There is no significant difference in exit rates between children of parents who own property and children of parents who do not. This result supports the claim that parents with no property are correctly interpreted as parents whose wealth is unobserved, not parents whose wealth is zero.

Table 4 reports results from different subsamples. Columns 1 and 2 show results for boys and girls separately. Parents' wealth has a significant effect on the probability of leaving the church only for boys, but the coefficients for boys and girls are close to each other in magnitude. Columns 3 and 4 show results for children of farmers and children

of non-farmers. The effect of parents' wealth on children's propensity to leave the church appears to be larger for the children of farmers. This is consistent with the idea that farmers bequeath a larger fraction of their wealth to their children, and so an increase in parents' wealth increases the effective punishment for leaving the church more for children of farmers than for children of non-farmers.

The broad result from tables 3 and 4 is that children of rich parents are less likely to leave the church than children of poor parents. There are two possible explanations for this finding. First, it may be the case that the return to physical or human capital investments is greater in the Amish community than in the surrounding society. In this case the Roy model predicts that rich households should prefer to remain Amish to take advantage of these greater returns, while poorer households should prefer to exit. Second, it may be the case that Amish households are not unitary, and that some Amish children prefer to leave the community while their parents prefer that they stay. In this case parents bribe their children to remain in the church by threatening to disinherit children who leave, and so children of rich parents are more likely to remain in the church, as in proposition 2.

Given the likelihood that Amish people do not realize the full potential return to their physical or human capital investments due to productivity restricting Amish rules, it is implausible that the return to investment in the Amish community is greater than the return to investment in the surrounding society. Therefore, I conclude from tables 3 and 4 that Amish households are not unitary, and that in at least some households the children prefer to leave the church while their parents prefer that they remain. However, by itself this result does not provide evidence for why a conflict between parents and children exists, and it does not provide evidence for whether and how the community plays a role in mediating intrahousehold conflict. For evidence on these questions I turn to the next step of the argument.

4.3 Step 2: Amish community punishments help to retain children in the church who would otherwise prefer to leave

The results of step 1 imply that there is conflict between Amish parents and their children. Amish parents want their children to remain in the church, while the children prefer to leave. Next I argue that Amish community rules, and in particular the punishment of shunning, help to retain children in the church who would otherwise prefer to leave. Thus Amish communities support parents in their conflict with their children.

I make this argument by comparing exit rates across different Amish affiliations. Figure 3 shows exit rates for the three affiliations covered by my dataset. The exit rate is highest for the New Order, intermediate for the Old Order, and lowest for the Andy Weaver affiliation. The differences between exit rates across affiliations are large. Recalling that the New Order are least likely to shun children who leave, the Old Order are intermediate, and the Andy

Weaver affiliation are most likely to shun, figure 3 provides strong evidence that shunning helps to reduce exit rates. Thus the community helps to support parents' desire to keep children in the church even when children prefer to leave.

In table 5 I report the coefficients on affiliation fixed effects in a regression of the child's probability of exit on affiliation, controlling for child's gender, child's age, the age of both parents, the father's occupation, the logarithm of the value of buildings owned by the parents, and an indicator for whether the parents own no property. Line 1 shows results from the full sample. The results confirm that even after controlling for observable differences between affiliations, there are large and statistically significant differences in exit rates between affiliations, consistent with the hypothesis that shunning helps to prevent children from exiting the church, as in proposition 3.

A different view of the relationship between religious rules and exit rates from the religious community comes from the club goods model. In the club goods model, part of the purpose of religious rules is to screen out members of the community who are less likely to contribute to the club good. Thus in religious groups described by the club goods model, religious rules increase the exit rate for at least some subset of the population. The screening function of religious rules is emphasized by Berman (2000) and Campante and Yanagizawa-Drott (2015). In order to test this mechanism I report exit rates by affiliation after controlling for observables for various subsamples in lines 2 through 7 of table 5. In every subsample, the exit rate is highest in the New Order, intermediate in the Old Order, and lowest in the New Order. Thus, contrary to the predictions of the club goods model, Amish rules do not serve to increase exit rates among any identifiable subset of the Amish population.

The results in table 5 also have implications for understanding cultural persistence across generations. Many authors have noted that minority cultures can often persist over time even in the face of material incentives to assimilate into the larger society. Two theories have been suggested in the economics literature for this persistence. Lazear (1999) and Michalopoulos (2012) argue that minority cultures persist because of technological barriers to assimilation. They argue that it is costly for speakers of minority languages to learn the majority language, and this cost prevents speakers of minority languages from adopting the majority culture. Bisin and Verdier (2000, 2001) argue that parents in minority cultures have a preference for following the practices of the minority culture, and that they instill a similar preference in their children. These unusual preferences ensure the persistence of the minority culture even in the face of material benefits from assimilation. The results in table 5 suggest that neither technological barriers nor parental preferences are sufficient to prevent substantial assimilation. In particular, without shunning the New Order affiliation is unable to prevent substantial exit. In contrast, strict shunning by the Andy Weaver affiliation is sufficient to reduce exit to very low levels. This suggests that a parental coalition aimed at

punishing children who assimilate can prevent assimilation when technological barriers and individual parental preferences cannot.

4.4 Step 3: Without the community, Amish parents have difficulty committing to punishing their children

So far I have argued that there is conflict between Amish parents and Amish children, and that by increasing the punishment for children who leave the church Amish communities help to support parents in this conflict. Next, I argue that without the support of the community parents have difficulty committing to punish their children.

The first piece of evidence for this claim is qualitative. Members of the Amish church in good standing are more likely to shun excommunicated family members in public than in private. Kraybill (2001, p. 139) describes this phenomenon as follows:

“The application of shunning varies widely from family to family. Many times it is relaxed in private homes but tightened in public settings if other church members are present, attesting to its ritual character and ceremonial role in the community. Many families treat family members under the ban [that is, family members who have been shunned] with love and care in the privacy of their homes. Despite its theological purposes, shunning is a painful process. One woman said, ‘I’m not responsible for being born into a church that practices shunning. I have an uncle and aunt and cousins in the ban, and it may separate us on a social level, but it could never sever the cord of love.’”

Similarly, Hurst and McConnell (2010, p. 92) write:

“In some cases family and church members are less likely to take the hard line on shunning when no one else is watching, as this ex-Swartzentruber male relates: ‘I have brothers that’ll shun somebody right in front of the whole family, but when they’re by themselves, they won’t.’”

In short, many Amish people do not want to shun their family members *ex post*, and avoid doing so when they are not monitored by the community.

I also provide quantitative evidence that Amish parents cannot commit to disinheriting their children without the support of the community. If Amish parents want their children to remain in the church so that their children will provide them with more care, then the parents may also want to prevent their children from migrating, again to increase the amount of care that the children provide. However, migration is not against the rules of the church, and so the church does not support parents who want to commit to disinheriting children who migrate. Thus I predict that while children of richer Amish parents are less likely to leave the church, there is no relationship between parent’s wealth and children’s propensity to migrate. Table 5 tests this hypothesis by running regressions analogous to the regressions in table 3, using an indicator for whether the child has migrated by 2010 as the dependent variable.

Column 1 of table 6 regresses an indicator for whether a child has migrated by 2010 on the measure of parent's wealth and an affiliation fixed effect. The effect of parents wealth on migration is not statistically significant, and the coefficient is positive. Column 2 adds control variables. The coefficient on parents wealth is qualitatively unchanged. Children of farmers are less likely to migrate, and boys are more likely to migrate. Column 3 adds each child's completed number of siblings as a control, which also does not qualitatively affect the results. Column 4 uses the alternate measure of parents' wealth, which also yields similar results.

Table 6 suggests that Amish parents' wealth has no effect on children's propensity to migrate. This result is consistent with the hypothesis that Amish parents commit to disinheriting children who leave the church, but cannot commit to disinheriting children who migrate, and so the results provide some support to my theory. However the results are also consistent with other hypotheses. In particular, parents' wealth may have a positive effect on children's propensity to migrate by alleviating borrowing constraints in addition to any negative effects due to the threat of disinheritance. Thus the results of table 5 should be treated with caution.

4.5 Step 4: Amish parents want their children to follow the rules of the Amish church because doing so strengthens family ties

Steps 1 through 3 of the argument establish that Amish parents and children are in conflict about whether the children should follow the rules of the Amish church, and that the Amish community supports parents in this conflict both by directly punishing children who leave and by helping parents to commit to punishing children who leave. However, so far the argument has not established why Amish parents want their children to follow the rules of the church. In this section I argue that the reason parents want children to follow the rules of the church is Amish church membership increases the quantity of resources that children devote to their parents and strengthens family ties.

As a preliminary, it is helpful to define the main alternative hypothesis for why Amish parents want their children to follow the rules of the church. The main alternative hypothesis is that Amish parents want their children to avoid modern technology and advanced education for some reason other than the effects of these choices on material outcomes. For example, Amish parents may believe that education and use of modern technology lead to morally corrupting contact with the outside world, that education and use of modern technology lead to punishment in the afterlife, or that education and technology are just intrinsically bad. Both my theory and these alternative theories are equally consistent with the evidence in steps 1-3.

In order to distinguish between these possibilities, I compare Amish institutions with the institutions of a closely related group, the Hutterites, whose society is described by Hostetler

(1974). The Hutterites and the Amish share very similar histories and theological beliefs. Like the Amish, the Hutterites are Anabaptists who migrated from Europe to the Americas in order to escape religious persecution, and like the Amish, the Hutterites continue to speak a dialect of German among themselves and to avoid many kinds of social contact with outsiders. Present day Hutterites live primarily in the prairie provinces of Canada, although some live in the northern plains states of the U.S. Hutterite theology, like Amish theology, emphasizes humility, simplicity, and surrender of the individual will to the will of God. In fact the Hutterites and the Amish use the same German word, *Gelassenheit*, to refer to the ideal state of spiritual and ethical virtue.

While the Hutterites and the Amish share similar histories and theological beliefs, they have different institutions and these institutions create different economic incentives. Most importantly, Amish households own property separately. In contrast, Hutterites live in agricultural communes called colonies, each with 70-150 members, and within the colony all property is owned collectively. Production decisions are made by the elders of the colony, and the output of the colony is shared equally among all members.

Because of their similar histories and theological beliefs, the Amish and the Hutterites presumably have similar preferences. Thus, if Amish rules reflect non-material preferences for avoiding technology and education, then the Hutterites should have similar preferences and hence similar rules. In contrast, if Amish rules are meant to solve material incentive problems, then the quite different incentive problems faced by the Hutterites suggest that the Hutterites might impose different rules even if they have similar underlying preferences. In particular, proposition 4 suggests that the Hutterites, with their much more extensive redistribution than the Amish, should have rules that are more hostile to family life.

Some of the rules imposed by the Amish and the Hutterites are similar. Both the Amish and the Hutterites have rules regarding dress and appearance, and these rules are nearly identical across the two societies. Both the Amish and the Hutterites also restrict the use of consumption technologies such as television. In other respects, though, Amish and Hutterite rules are very different. The Hutterites place no restrictions on the use of productive technology. Members of Hutterite colonies drive cars and use modern agricultural technology, including computerized and internet enabled devices for managing various aspects of farm production. There are restrictions on education in the sense that decisions about whether colony members are allowed to pursue higher education must be made by the entire colony. However, many Hutterites graduate from high school, and some attend university. On the other hand, the Hutterites place significant restrictions on family life that do not exist in Amish society. Hutterite children are separated from their parents at the age of 2 and spend most of their time in the colony kindergarten, being supervised and instructed by colony teachers. All Hutterite meals are taken communally in the colony dining hall, and at meals men, women, and children sit in separate sections, so that husbands do not eat with wives

and parents do not eat with children. Hutterite children do share rooms for sleeping and storing possessions with their parents, within larger apartment complexes on the Hutterite colony property. However, the doors on these rooms typically have no locks and Hutterites enter each other's rooms without knocking, so that there is very little private family space. Table 7 summarizes the differences between Amish and Hutterite rules.

The differences between Hutterite and Amish rules are consistent with the hypothesis that Hutterite rules are meant to encourage people to contribute effort to communal production, while Amish rules are meant to encourage people to spend more time with their families. Time spent on private consumption, such as time spent watching television, subtracts both from time devoted to family and from time devoted to the community, and so both the Hutterites and the Amish restrict consumption technologies. However, Amish and Hutterite attitudes towards productive technologies differ. For the Amish, time spent on market production subtracts from time spent with family members, and so the Amish impose restrictions on productive technology to increase time spent with family members. In contrast, for the Hutterites effort contributed to communal production benefits the entire community, and so it would make no sense for the Hutterite community to impose restrictions lowering the productivity of work effort. On the other hand, for the Hutterites time devoted to family subtracts from time devoted to communal production, and so it does make sense for the Hutterites to impose rules restricting time spent with families.

Given that the Hutterites and the Amish have very different rules regarding family life, it would be reasonable to expect that Hutterite and Amish family relationships would differ, and in fact this is the case. Hutterite families are not very close. Hostetler (1974, p. 203) describes the relationship between the Hutterite colony and the Hutterite family as follows:

“The Hutterite colony functions in many ways like an extended family. Because Hutterite society has institutionalized a continuing relationship between parents and children, the family is emotionally less demanding and less exclusive than is the rule with middle-class Americans.”

In contrast, the Amish believe that the rules of the church strengthen their family ties. Kraybill and Nolt (1995, p. 126-127) explain Amish church rules preventing businesses from growing too large as follows:

“[The Amish] worry that bigger establishments will weaken the involvement of the family. One of the reasons they established small shops was to keep work integrated with family life. A large business easily encroaches on family life, as entrepreneurs become engulfed in manufacturing, sales, or bookkeeping. Although the community values hard work, it frowns on jobs that completely dominate everything else. ‘A lot of Amish are workaholics,’ said an Amishman who manufactures storage barns. ‘They pretty easily get too caught up in their work. The church doesn’t want a business to get too awfully big, or pretty soon you’re living for your job. You get too caught up in it. That’s not right.’”

Amish families are in fact very closely knit. Kraybill (1989, p. 88) describes the importance of the Amish family as follows:

“The [Amish] family’s scope and influence dwarfs that of the modern nuclear family. Amish life is spend in the context of the family. In contemporary families social functions from birth to death, from eating to leisure, often leave the home. In contrast, Amish activities are anchored at home. Children are usually born there. They play at home and walk to school. By age fourteen, children work full-time in the home, shop, or farm. They are taught by their extended family, not by television, babysitters, popular magazines, or daycare teachers. Young couples are married at home. Church services rotate from home to home. Most meals are eaten at home. Adults work at home or nearby.”

To summarize, the Amish and Hutterite rules differ in ways that are consistent with the hypothesis that the rules of the two societies are designed to affect material outcomes, and not consistent with the hypothesis that the Amish and the Hutterites have intrinsic preferences for following rules that are unrelated to their material effects. In particular, Amish rules are designed to increase the amount of time that family members devote to each other, while Hutterite rules are designed to reduce the amount of time that family members devote to each other and to increase the amount of time that family members devote to the community.

5 Further evidence

I have argued that strict religious groups can follow one of two patterns. Right-wing religious groups feature relatively less redistribution, attract rich members, and have rules that increase the quantity of resources that family members devote to each other. Left-wing religious groups feature relatively more redistribution, attract poor members, and have rules that are neutral towards the family or that may even reduce the quantity of resources that family members devote to each other. The Amish are an example of a right-wing religious group, while the Hutterites are an example of a left-wing religious group. In this section, I discuss some other examples of this dichotomy.

5.1 Quakers and Shakers

The Quakers are a Protestant sect that appeared in England in the 17th century. Many Quakers subsequently migrated to the Americas, and in particular to Pennsylvania, in part to escape persecution in England. The Shakers branched off from the Quakers in England in the 18th century and subsequently also migrated to the Americas, primarily to upstate New York and other parts of New England.

Quakers in colonial Pennsylvania owned private property, and Quaker society was both relatively wealthy and relatively unequal. Table 8 shows data from the Philadelphia tax lists

of 1754 and 1774, for the whole population and for Quakers. Each entry is the percentage of households in each wealth category. Quakers appear to be richer than the overall population, and there is also significant inequality within the Quaker community.

Quakers imposed many rules on their members, enforced through regular meetings at which deviant members could be disciplined or possibly expelled from the church. Rules emphasized “plain living” and prohibited many kinds of ostentatious consumption, drinking, and gambling. Frost (1973, p. 56) lists the offenses that could be punished at Quaker meetings, including “marriages without Friends’ [that is, Quakers] approval or by an Anglican or Presbyterian clergyman, sexual offenses, drunkenness, gossip or slander, military service or privateering, quarrelling or fighting, bankruptcy, going to law against a fellow Quaker, dishonesty, profanity, swearing, playing cards, abusive behavior to wife or children, and attendance at plays or horse races.”

One important purpose of these rules was to control children’s behavior. Frost (1973, p. 134) describes the process as follows:

“If parents found it impossible to control their children, they could always turn to the meetings for help. The process was similar to that used in disciplining. At first a few Friends would talk to the child privately. If he remained obstinate, the matter was presented to the monthly meeting and an official delegation counseled him. If the youth still did not reform, then the meeting took steps to disown him.”

Without the church rules, many Quaker parents had difficulty committing to disciplining their children. Frost (1973, p. 134) writes:

“The necessity of repeated exhortations [by the Quaker community] to parents to inculcate in their children the value of plainness in speech and apparel and to restrain them from mixing with the world must mean that some mothers and fathers slighted these duties.”

Frost (1973, p. 135) also recounts the following anecdote:

“James Logan [a Pennsylvania Quaker] once threatened disinheritance if his son did not write home more often (Logan had received one letter in a year), but the threat was mostly bluff, and Logan soon sent assurances that the boy would be well treated on his return.”

Possibly as a result of the community rules, Quaker families were tightly knit. Quakers emphasized respect and companionship between husbands and wives. Marriage, according to the Quakers, was a “blessed state entered into so that husband and wife might fulfill each other intellectually, spiritually and physically.” (Frost 1973, p. 150). Children were highly valued by their parents.

The Shakers held theological beliefs that were similar to the Quakers but had a very different form of economic organization. Like the Hutterites, the Shakers lived in agricultural communes in which all property was owned jointly and the output of the commune was shared equally. One problem that arose from this practice was that poor outsiders would attempt to join Shaker communities for the economic benefits, and not because they agreed

with Shaker theology. Shakers described these free-riders as “winter Shakers” or “‘bread and butter’ Believers”, who were “interested only in such short-term physical benefits, as food, shelter, and clothing.” (Stein 1992, p. 162).

Shaker attitudes towards the family were also very different from Quaker attitudes. Shakers practiced celibacy, and as a result members of Shaker communities did not have spouses or biological children. Shaker communities sustained themselves instead by adopting runaways and abandoned babies, and raising them communally. Shaker communities acted in many ways as surrogate families, and in fact Shakers referred to their communities as “families” and addressed each other as “brother” or “sister”.

To summarize the Quakers appear to be an example of a right-wing religious group in my classification, while the Shakers are an example of a left-wing religious group. Quaker society was relatively unequal and attracted the wealthy, and Quaker rules supported the family. In contrast Shaker society was perfectly equal and attracted the poor, and Shaker rules prevented the formation of biological families and encouraged people to contribute instead to the community as a kind of surrogate family.

5.2 Sadducees, Pharisees, and Essenes

The Sadducees, Pharisees, and Essenes were Jewish sects that were prominent in Roman Palestine during the Second Temple period, as described by the ancient historians Josephus and Philo. The Sadducees and the Pharisees lived in Jerusalem and competed for political control over the Temple and by extension over much of Jewish society. The Sadducees represented the wealthy, while the Pharisees represented both wealthy and poorer classes (Newman 2006, p. 58 and 76). Both the Sadducees and the Pharisees placed great importance on following the Jewish law, including rules forbidding many kinds of consumption. A central function of the law was to help parents control children, as exemplified by the fifth commandment, “Honor thy father and thy mother.”

In contrast to the Sadducees and the Pharisees, the Essenes lived outside of Jerusalem in rural communes. Like the Hutterites and Shakers, they owned property communally (Newman 2006, p. 86). Also like the Shakers, they practiced celibacy and maintained the community by adopting other people’s children (Newman 2006, p. 87). Philo explains Essene celibacy as follows:

“Furthermore, they eschew marriage because they clearly discern it to be the sole or the principal danger to the maintenance of communal life... For he who is either fast bound in the love lures of his wife or under the stress of nature makes his children his first care ceases to be the same to others.” (Newman 2006, p. 166)

Thus the Pharisees and Sadducees appear to be right-wing religious groups in my classification, while the Essenes are a left-wing religious group.

6 Conclusion

In this paper I have constructed a model of religion as an institution that provides community enforcement of contracts within families. I have then provided evidence that this mechanism helps to explain the organization of Amish religious communities. By way of conclusion, I reiterate the main generalizable results from the paper, and discuss some policy implications:

1. Religious groups differ in the values that they emphasize. Right-wing religious groups feature relatively less redistribution and have rules that strengthen family ties. These groups attract richer members. Left-wing religious groups feature relatively more redistribution and have rules that are neutral towards the family or that may even weaken family ties. These groups attract poorer members. The differences between these groups are due to the different incentive problems faced by their members and not to differences in fundamental preferences. Right-wing groups are described by my model while left-wing groups are described by the club goods model of Iannaccone (1992).
2. Religious groups described by my model are inefficient, in the sense that it would be possible for parents and children to achieve a Pareto superior outcome if parents could commit to bequest schedules conditional on their children's actions. The source of the inefficiency is that not only punishment but also monitoring of children must be delegated from the parents to the community, but the community has less information than parents about children's actions. Seemingly inefficient cultural behaviors really are inefficient in my theory, in contrast to the theories of Boyd and Richerson (1985) and Bisin and Verdier (2000, 2001) in which seemingly inefficient cultural behaviors are in fact efficient given people's preferences and cognitive abilities.
3. Despite the inefficiency of religious groups described by my model, these groups can persist over many generations. Neither language barriers, as in Lazear (1999) and Michalopoulos (2012), nor parental preferences, as in Bisin and Verdier (2000, 2001) are sufficient by themselves to maintain cultural distinctiveness in the face of material incentives for assimilation into the larger society. However a coalition of parents who collectively agree to punish children who assimilate can successfully maintain cultural distinctiveness over long periods of time.
4. Since minority cultural persistence is driven by parents' desire to get care from their children, and not by technological barriers to assimilation or by intrinsic preferences for engaging in distinctively minority behaviors, changes in public policy may affect whether minority cultures persist. In particular policies that provide old age support, or that otherwise substitute for services normally provided by the family, may increase the rate of cultural assimilation. However, it is not clear whether such policies are

desirable from a welfare perspective. The value of family togetherness and the value of economic efficiency are not fully commensurable, and decisions between these values cannot be made exclusively on the basis of revealed preference analysis. Thinking about how to trade off these values would be an interesting project for research in both economics and philosophy.

References

- [1] Abramitzky, Ran (2008), “The Limits of Equality: Insights from the Israeli Kibbutz,” *Quarterly Journal of Economics*, 123:3, 1111-1159.
- [2] Akerlof, George A. and Rachel E. Kranton (2000), “Economics and Identity,” *Quarterly Journal of Economics* 115:3, 715:753.
- [3] Behrman, Jere R. and Mark R. Rosenzweig, “Parental Wealth and Adult Childrens Welfare in Marriage,” *Review of Economics and Statistics* 88:3, 496-509
- [4] Berman, Eli (2000), “Sect, Subsidy and Sacrifice: An Economist’s View of Ultra-Orthodox Jews,” *Quarterly Journal of Economics* 115:3, 905-953.
- [5] Bernheim, B. Douglas, Andrei Shleifer, and Lawrence H. Summers (1985), “The Strategic Bequest Motive,” *Journal of Political Economy* 93:6, 1045-1076.
- [6] Bisin, Alberto and Thierry Verdier (2000), “Beyond the Melting Pot: Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits,” *Quarterly Journal of Economics* 115:3, 955-988.
- [7] Bisin, Alberto and Thierry Verdier (2001), “The Economics of Cultural Transmission and the Dynamics of Preferences,” *Journal of Economic Theory* 97:2, 298-319.
- [8] Botticini, Maristella and Zvi Eckstein (2012), *The Chosen Few: How Education Shaped Jewish History, 70-1492*, Princeton: Princeton University Press.
- [9] Bowen, T. Renee, David M. Kreps, and Andrzej Skrzypacz (2013), “Rules with Discretion and Local Information,” *Quarterly Journal of Economics* 128:3, 1273-1320
- [10] Buser, Thomas (2015), “The Effect of Income on Religiousness,” *American Economic Journal: Applied Economics* 7:3, 178-195.
- [11] Chen, Daniel L. (2010), “Club Goods and Group Identity: Evidence from the Islamic Resurgence during the Indonesian Financial Crisis”, *Journal of Political Economy* 118:2, 300-354.
- [12] Dehejia, Rajeev, Thomas DeLeire, and Erzo F.P. Luttmer (2007), “Insuring consumption and happiness through religious organizations,” *Journal of Public Economics* 91:1-2, 259-279.
- [13] Frost, J. William (1973), *The Quaker Family in Colonial America: A Portrait of the Society of Friends*, New York: St. Martin’s Press.

- [14] Ghosh, Suman and Alexander Karaivanov (2007), “Can a raise in your wage make you worse off? A public goods perspective” *Journal of Development Economics* 84:1, 551-571.
- [15] Greif, Avner (1993), “Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Trader’s Coalition,” *American Economic Review* 83:3, 525-548.
- [16] Hostetler, John A. (1993), *Amish Society*, Baltimore: Johns Hopkins University Press.
- [17] Hostetler, John A. (1997), *Hutterite Society*, Baltimore: Johns Hopkins University Press.
- [18] Hurst, Charles E. and David L. McConnell (2010), *An Amish Paradox*, Baltimore: Johns Hopkins University Press.
- [19] Iannacone, Lawrence (1992), “Sacrifice and Stigma: Reducing Free-Riding in Cults, Communes, and other Collectives,” *Journal of Political Economy* 100:2, 271-291.
- [20] Kandori, Michihiro (1992), “Social Norms and Community Enforcement,” *Review of Economic Studies* 59:1, 63-80.
- [21] Kraybill, Donald B. (2001), *The Riddle of Amish Culture*, Baltimore: Johns Hopkins University Press.
- [22] Kraybill, Donald B., Karen M. Johnson-Weiner, and Steven M. Nolt (2013), *The Amish*, Baltimore: Johns Hopkins University Press.
- [23] Kraybill, Donald B. and Steven M. Nolt (2004), *Amish Enterprise*, Baltimore: Johns Hopkins University Press.
- [24] Lazear, Edward P. (1999), “Culture and Language”, *Journal of Political Economy* 107:S6, S95-S126.
- [25] Michalopoulos, Stelios (2012) “The Origins of Ethnolinguistic Diversity”, *American Economic Review* 102:4, 1508-1539.
- [26] Munshi, Kaivan and Mark R. Rosenzweig, ”Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap,” *American Economic Review*, forthcoming.
- [27] Newman, Hillel (2006), *Proximity to Power and Jewish Sectarian Groups of the Ancient Period*, Leiden: Koninklijke Brill.
- [28] Stein, Stephen J. (1992), *The Shaker Experience in America*, New Haven: Yale University Press.

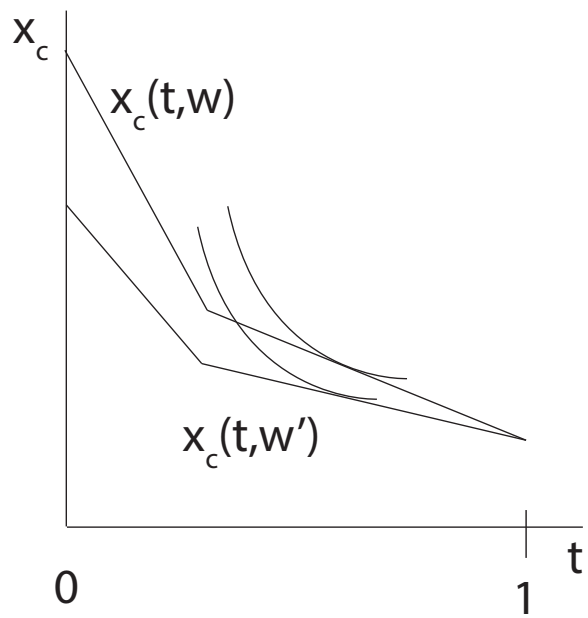


Figure 1: This figure shows budget sets and utility-maximizing indifference curves for a child with wage w and w' , where $w' < w$. The child always gets greater utility from the higher wage.

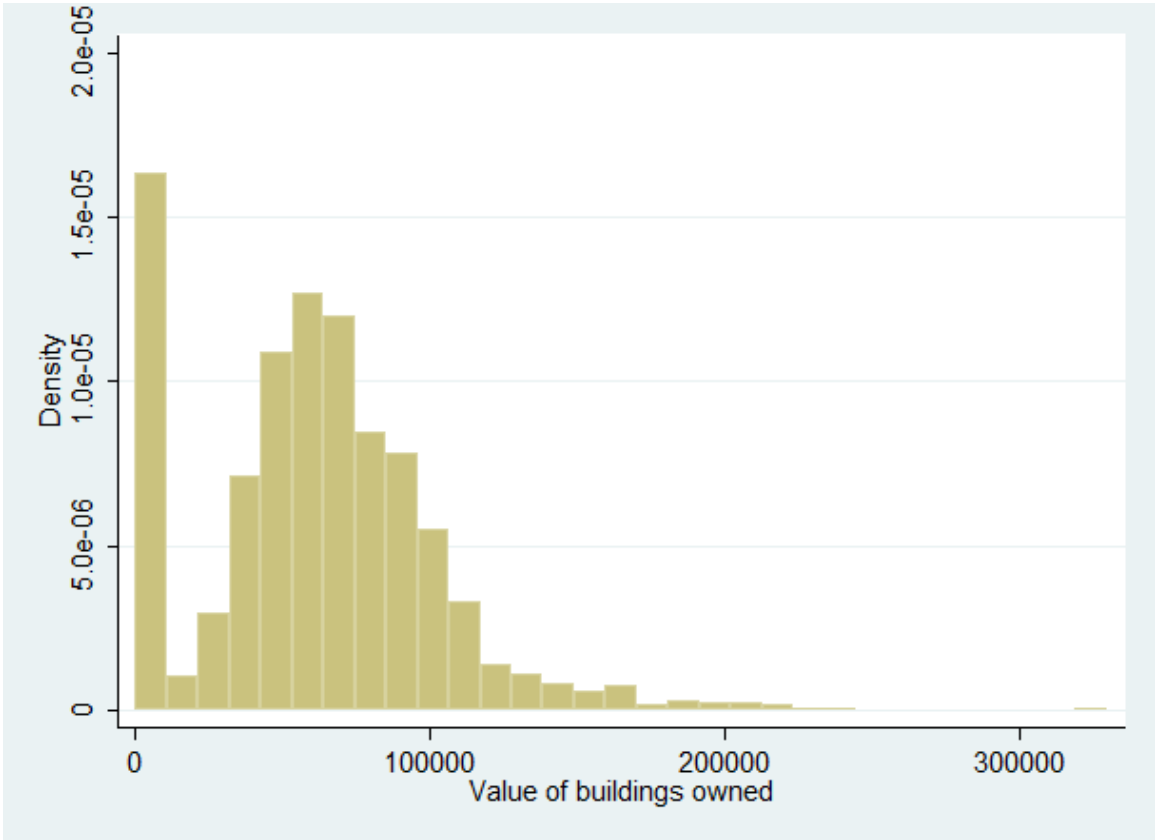


Figure 2: Wealth distribution of parents in sample

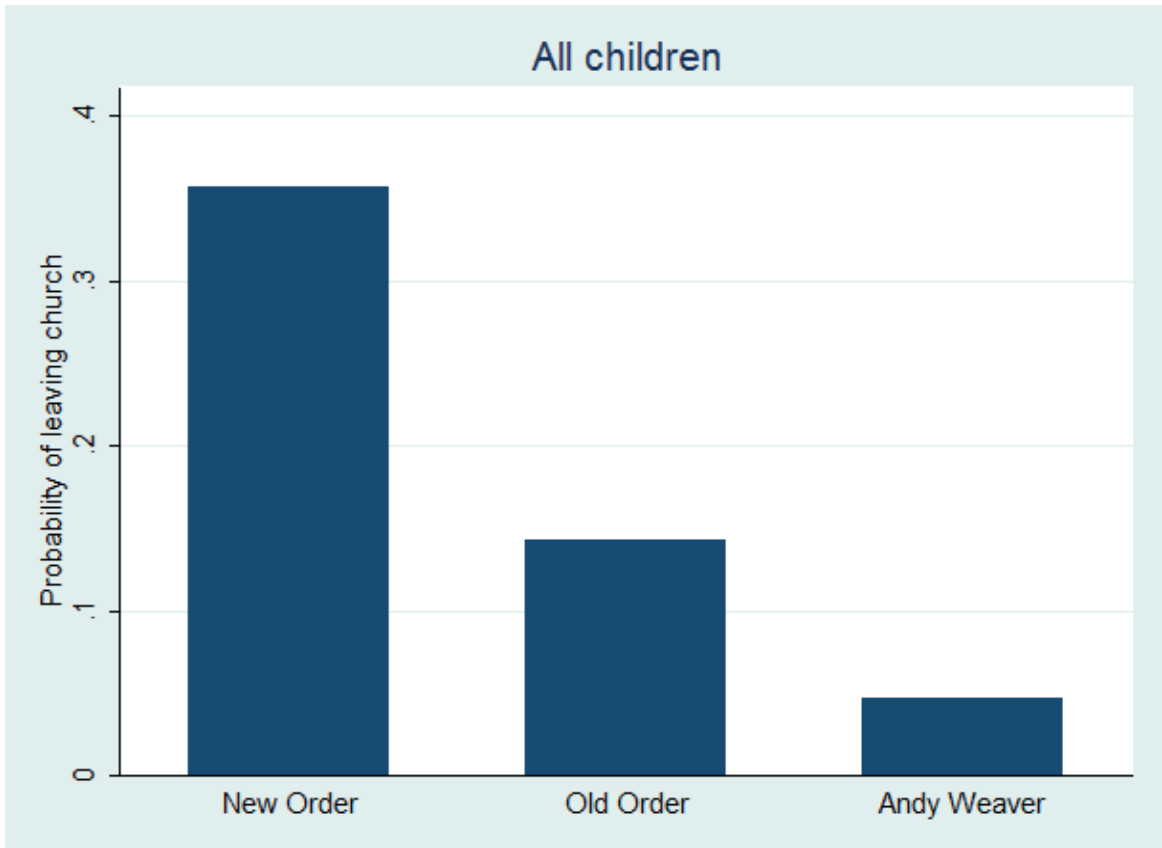


Figure 3: Probability of leaving church by affiliation for all children in sample

Table 1: North American Amish population estimates

Year	Estimated population
1901	6,300
1911	6,600
1921	9,960
1931	14,300
1941	21,100
1951	27,675
1961	36,855
1971	50,280
1981	77,955
1991	123,025
2001	189,335
2012	273,710

Source: Kraybill, Johnson-Weiner and Nolt (2013), p. 156

Table 2: Summary Statistics

	All matched observations	Observations with building value>0
Outcome variables		
Child leaves church	0.146 (0.353)	0.147 (0.354)
Child migrates	0.0656 (0.248)	0.0626 (0.242)
Child-level variables		
Child's age	11.79 (2.579)	11.90 (2.580)
Child is male	0.511 (0.500)	0.508 (0.500)
Household-level variables		
Father is farmer	0.513 (0.500)	0.493 (0.500)
Father's age	41.10 (6.390)	41.62 (6.416)
Mother's age	40.07 (6.337)	40.57 (6.330)
Number of siblings	7.836 (3.109)	7.835 (3.150)
Building value	59,803 (41,112)	72,657 (33,467)
Observations	3,611	2,972

Standard errors in parentheses. Dollar values are 1988 dollars.

Table 3: Effect of parents' wealth on children's decision to leave the Amish church

Mean of dependent variable	0.147	0.147	0.147	0.147	0.145	0.146
log(Parent's wealth)	-0.0516*** (0.0174)	-0.0458** (0.0183)	-0.0458** (0.0178)			
log(Parent's wealth) (alternate measure)					-0.0626** (0.0278)	
Parents own no property						0.0193 (0.0214)
Child is male		0.0573*** (0.0147)	0.0573*** (0.0147)	0.0573*** (0.0147)	0.0544*** (0.0151)	0.0545*** (0.0135)
Child's age		0.00145 (0.00255)	0.000847 (0.00250)	0.000847 (0.00250)	0.00189 (0.00311)	0.00304 (0.00224)
Father is farmer		-0.0778*** (0.0188)	-0.0681*** (0.0203)	-0.0681*** (0.0203)	-0.0561*** (0.0310)	-0.0916*** (0.0174)
Father's age		0.0106*** (0.00348)	0.0116*** (0.00338)	0.0116*** (0.00338)	0.0142*** (0.00398)	0.00826*** (0.00305)
Mother's age		-0.00653* (0.00331)	-0.00692** (0.00327)	-0.00692** (0.00327)	-0.0104** (0.00398)	-0.00581** (0.00297)
Completed number of siblings						
Affiliation FE?	YES	YES	YES	YES	YES	YES
Observations	2972	2972	2972	2972	2120	3611

Dependent variable=1 if the child has left the Amish church by 2010. Wealth is the value of buildings owned by the child's parents in 1988. Wealth (alternate measure) is the value of buildings and land owned by the child's parents in 1988. No property=1 if the child's parents owned no buildings in 1988. Columns 1, 2, 3, and 4 drop observations of children whose parents own no buildings in 1988. Column 4 drops observations from Wayne and Tuscarawas counties. Standard errors clustered by district in parentheses. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 4: Heterogeneous effects of parents' wealth on children's decision to leave the Amish church

Sample	Boys	Girls	Father is farmer	Father is non-farmer
Mean of dependent variable	0.175	0.119	0.103	0.190
log(Parent's wealth)	-0.0584** (0.0249)	-0.0313 (0.0201)	-0.0657** (0.0267)	-0.0261 (0.0284)
Child is male			0.0381* (0.0200)	0.0750*** (0.0201)
Child's age	0.00507 (0.00384)	-0.00291 (0.00376)	0.00133 (0.00312)	0.00108 (0.00415)
Father is farmer	-0.0942*** (0.0277)	-0.0603*** (0.0183)		
Father's age	0.00496 (0.00459)	0.0166*** (0.00418)	0.00634 (0.00473)	0.0153*** (0.00513)
Mother's age	-0.00114 (0.00453)	-0.0120*** (0.00413)	-0.00257 (0.00395)	-0.0111** (0.00485)
Observations	1,511	1,461	1,466	1,506

Dependent variable=1 if the child has left the Amish church by 2010. Wealth is the value of buildings owned by the child's parents in 1988. All regressions drop observations for which parents own no buildings. Standard errors clustered by district in parentheses. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 5: Exit rates by affiliation with controls in full sample and various subsamples
 Dependent variable=1 if child leaves church by 2010

Sample	Coefficient on New Order	Coefficient on Old Order	Observations
Full sample	0.286 (0.0371)	0.0777 (0.0162)	3611
Boys	0.310 (0.0518)	0.0951 (0.0213)	1845
Girls	0.257 (0.0430)	0.0589 (0.0194)	1766
Father is farmer	0.255 (0.0492)	0.0516 (0.0178)	1854
Father is non-farmer	0.335 (0.0569)	0.114 (0.0337)	1757
Parents' wealth above median	0.252 (0.0463)	0.0833 (0.0194)	1488
Parents' wealth below median	0.376 (0.0564)	0.111 (0.0277)	1484
Parents own no property	0.207 (0.106)	0.003 (0.052)	639

Each line reports coefficients on affiliation fixed effects for a regression of an indicator for whether the child leaves the church on affiliation fixed effects plus controls. Controls include the logarithm of the value of buildings owned by the child's parents, an indicator for whether the child's parents own no buildings, an indicator for whether the child's father is a farmer, the child's age and gender, and the ages of the child's mother and father. The omitted fixed effect is for the Andy Weaver affiliation. Standard errors in parentheses are clustered by district.

Table 6: Effect of parents' wealth on children's decision to migrate

Mean of dependent variable	0.063	.063	0.063	0.055
log(Parents' wealth)	0.00303 (0.00989)	0.00522 (0.0103)	0.00520 (0.0104)	
log(Parents' wealth)(alternate measure)				0.00286 (0.0138)
Child is male		0.0173* (0.00884)	0.0172* (0.00884)	0.00623 (0.0104)
Child's age		-0.000537 (0.00203)	-0.000209 (0.00207)	-0.000826 (0.00220)
Father is farmer		-0.0235** (0.0104)	-0.0287** (0.0110)	-0.0158 (0.0154)
Father's age		0.000289 (0.00242)	-0.000263 (0.00245)	0.00258 (0.00272)
Mother's age		0.00102 (0.00231)	0.00124 (0.00232)	-0.00202 (0.00258)
Completed number of siblings			0.00377* (0.00218)	
Affiliation FE?	YES	YES	YES	YES
Observations	2,972	2,972	2,972	2,120

Dependent variable=1 if the child has migrated by 2010. Wealth is the value of buildings owned by the child's parents in 1988. Wealth (alternate measure) is the value of buildings and land owned by the child's parents in 1988. Columns 1, 2, 3, and 4 drop observations of children whose parents own no buildings in 1988. Column 4 drops observations from Wayne and Tuscarawas counties. Standard errors clustered by district in parentheses. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 7: Comparing Amish and Hutterite Institutions

	Amish (Private property)	Hutterites (Communal property)
Restricts consumption choices	Yes	Yes
Restricts productive technologies	Yes	No
Restricts education	Yes	No
Restricts family contact	No	Yes
Family ties	Strong	Weak

This table summarizes the differences between Amish and Hutterite rules and family ties. See text for details.

Table 8: Distribution of wealth of Quakers and non-Quakers in colonial Philadelphia

Assessed wealth in 1754 (£)	Total Population	Quakers
10-19	50	34
20-29	15	15
30-39	21	22
40-49	7	7
50-59	3	4
60-79	3	5
80-99	2	6
100-	4	6

Assessed wealth in 1774 (£)	Total Population	Quakers
0	53	26
1-9	17	17
10-19	9	8
20-29	5	6
30-39	4	6
40-49	2	2
50-59	2	6
60-79	3	6
80-99	2	6
100-149	3	5
150-199	1	2
200-	3	8

This table shows the distribution of assessed wealth for the total population of Philadelphia and for Quakers, derived from the Philadelphia tax list. Each entry is a percentage of the relevant population. Source: Frost (1973), p. 205